Exposure and amplitude limits for eradicating populations with inherited cellular states

20 September 2026

Abstract

How much intervention is necessary to eradicate a finite population whose cells inherit a molecular state, when the intervention may be chosen using the entire population history? We prove two independent obstructions for finite-type branching populations under bounded predictable feedback. The first is an exposure floor: a baseline extinction supersolution together with a componentwise bound on the intervention generator produces a family of supersolutions indexed by the remaining budget, and the resulting population barrier retains founder-count dependence without assuming that families are independent under feedback; for n identical founders it gives survival probability at least $1-(1-\eta e^{-cB})^n$. The second is an amplitude floor: a single supersolution valid across the whole admitted action range bounds survival away from zero for *every* policy, at every budget and every horizon. Necessary and sufficient exposures match at order $\log(n/\delta)$ above the amplitude threshold, and an outward-rounded Taylor calculation certifies a finite-time policy at exposure 11.6 with risk below 0.01 from one founder of the two-site molecular source. We then locate the amplitude threshold exactly in that source as a function of memory size: the critical intrinsic erasure is bracketed to six decimals for $2 \le N \le 8$ sites and rises from .0927 to .4506, so at five to eight sites no policy of amplitude .29 eradicates, whatever its budget or duration; explicit rational supersolutions bound the surviving probability below by .17, .46, .61 and .70 at $N = 5, \dots, 8$. A second actuator repairs this: added death on protected states has an N-uniform threshold $b - d_{\mathrm{prot}}$, gives a dimension-free sufficient exposure $1.45\log(n/\delta)$, and attains the exposure floor up to a factor below 1.54. We also record a preparation-dependent failure of monotonicity in erasure, and conditional concentration and administered-amount lower bounds. Probability proofs are conventional; source certificates and the finite-time enclosure are exact finite computations; selected algebra is checked in Lean.

Contents

1 Introduction

An inherited state can affect both the descendants a cell produces and their susceptibility to killing. Eliminating a marker along one lineage therefore need not eliminate the family. Conversely, a declining expected population can give a useful extinction guarantee, but a negative mean growth rate alone does not quantify finite-founder survival. The appropriate control question concerns the probability of leaving any surviving family under an explicitly priced intervention.

Reaction-based models connect histone and DNA modifications to cellular memory [4, 5, 11], and inherited non-genetic states are a standard explanation for drug-tolerant persistence [2, 16–18].

Inheritance and treatment resistance have been studied in population models [7]; drug-induced plasticity changes optimal dosing, and the chosen dose–response law can change which schedules are favored [12]. Controlled branching processes have a substantial stochastic-control theory [8, 13], and tumour elimination probabilities have been used directly in treatment objectives [6]. We do not claim priority for branching generating functions, stochastic verification, or extinction-based treatment objectives. Our contribution is a pair of explicit, checkable certificates — one indexed by the remaining budget, one uniform over the admitted action range — their founder-count and memory-size consequences, and a source-level realization that retains correlated molecular inheritance.

Two obstructions, and a third

Write B for the total intervention exposure, meaning additional molecular reaction rate integrated over time, and $v_{\max}$ for the largest instantaneous rate the actuator can deliver. Three separate quantities can prevent eradication.

- **Exposure.** Theorem 3.3 shows that if a baseline extinction supersolution q obeys a componentwise inequality $Rq \le c(1-q)$ for the intervention generator, then survival from configuration z is at least $1-\prod_i[1-e^{-cB}(1-q_i)]^{z_i}$ for *every* predictable policy respecting the pathwise budget B. The exposure needed to reach risk δ from n founders therefore grows like $c^{-1}\log(n/\delta)$.
- **Amplitude.** Theorem 4.1 shows that if a single q is a supersolution for every admitted action — equivalently, by affineness, at both ends of the action range — then survival is at least $1-\prod_i q_i^{z_i}$ for every policy at *every* budget and horizon. This obstruction is not removed by treating longer or paying more.
- **Time.** Proposition 4.4 shows that bounded death rates impose a minimum duration for finite-time extinction, independently of the two bounds above.

The three are logically independent, and they fail in different ways. Only the first is a budget statement; only the second survives an unbounded course; only the third constrains a fixed deadline.

What is new relative to the exposure floor alone

The central technical device for the exposure floor is that no optimization over a discretized population space is required. Given a baseline supersolution q, one considers the entire segment $1-a(1-q)$, which is preserved by the inequality a population barrier needs (Lemma 3.1); setting a by the remaining budget gives a bound that approaches one with founder count at every fixed finite exposure. A constant admissible action with a negative weighted drift gives the matching logarithmic order (Corollary 5.2), and an outward-rounded Taylor enclosure certifies a sharper finite-horizon policy in the worked source (Section 6.2).

The amplitude floor is what makes the memory-size question answerable. In the molecular source of Section 6 the erasure actuator has a critical amplitude $v_c(N)$: below it the certificates exhibited here show that *no* policy eradicates, and above it a constant action does. Section 7 brackets the corresponding critical intrinsic erasure to six decimals for $2 \le N \le 8$ sites; it rises from .092750 at two sites to .450572 at eight, a factor of 4.86. In particular the amplitude .29 used in the two-site worked example is certified *insufficient at any budget or duration* for $5 \le N \le 8$, with explicit rational supersolutions bounding survival from a fully marked founder below by .170, .460, .613 and .701 at $N = 5, 6, 7, 8$. This settles, negatively, a question left open by a coarser mean-growth diagnostic, which had only shown that a particular constant action ceases to contract.

Section 7 also shows that the deterioration is a property of this actuator and not of inherited memory as such. A second actuator that adds death on protected states has threshold at most

$b - d_{\text{prot}}$ for every N, because beyond that value the expected total count decays at rate $\min(\kappa - (b - d_{\text{prot}}),\, d_{\text{unprot}} - b)$ with unit weight vector; the resulting sufficient exposure $1.45 \log(n/\delta)$ is dimension-free and lies within a factor 1.54 of the exposure floor at $n = 1$, $\delta = .01$. Where the eraser alone needs a fourfold amplitude increase between two and eight sites, a fixed added death of .025 restores mean contraction at six sites at the erasure amplitude that by itself provably fails.

Scope

The case study uses synthetic reaction and demographic rates. It is suitable for exposing mathematical dependencies and for designing a low-density microcolony test, not for prescribing a drug regimen. The exposure variable measures additional molecular reaction rate integrated over time; it is not a drug mass or a concentration, and Section 8.2 states exactly what is needed to convert it. Throughout, "eradication" means the specified extinction event in a process with no immigration, and "survival" means non-extinction of the mathematical population, which is not a detection threshold, a clinical recurrence or a progression endpoint.

2 Controlled inherited-state populations

Let $\{1, \dots, m\}$ be a finite type space and $Z_t \in \mathbb{N}^m$ the count vector. A cell of type i changes type according to a conservative generator Q, dies at rate d_i, and is replaced at rate b_i by an offspring vector with probability generating function $F_i(x)$. Offspring types may be correlated; the joint offspring law is part of the data. We assume a deterministic upper bound K on offspring count, and that all per-cell rates are bounded over the admitted action set. Define

$$\Phi(x) = Qx + d \odot (1 - x) + b \odot (F(x) - x), \tag{1}$$

where $\odot$ is the entrywise product. For binary offspring write $F_i(x) = D_i(x, x)$ with D_i the bilinear form of the ordered daughter law; the mother's removal accounts for the term $-b_i x_i$.

Under deterministic rates, conditional independence of future offspring clocks given their birth states gives the backward equation $\dot{u} = \Phi(u)$. With terminal datum zero its solution is extinction by the horizon. A terminal datum q^{off} instead represents eventual extinction under a specified continuation after the course. For piecewise constant schedules the autonomous flows act on this terminal datum in reverse chronological order.

2.1 Actions, feedback and budget

During intervention use $Q_v = Q_{\text{base}} + vR$ with $R\mathbf{1} = 0$ and $0 \le v_t \le v_{\max}$, each Q_v a valid generator. We write Φ_{base} for (1) at $v = 0$ and Φ_v for the field with Q_v in place of Q. The scalar common action is predictable with respect to the entire population history and any independent controller randomization. It obeys the pathwise constraint

$$C_t = \int_0^t v_s \, ds \le B \qquad (0 \le t \le T) \qquad \text{almost surely.} \tag{2}$$

The constraint is pathwise, not an expectation constraint: replacing (2) by $\mathbb{E}C_T \le B$ requires a different argument. Nor is B a sum of cell-specific exposures; the common-control budget is a single number spent by a single actuator. The background death and reproduction rates are fixed unless an extension explicitly changes them. This full-information feedback class contains any controller using less information. Families need not be independent under it.

The process is constructed by predictable thinning of bounded Poisson proposals on a countable genealogical tree. Birth count is dominated by a bounded-offspring pure-birth process, with expected size at most $|Z_0|e^{b_{\max}(K-1)t}$ for $K \geq 1$. Expected proposal counts on finite intervals are finite, which proves non-explosion by localization. Empty populations remain empty, so 0 is absorbing and there is no immigration. Feedback changes conditional intensities; it does not invalidate this construction. A *continuation* is a specified uncontrolled branching source, started with fresh randomness conditional on the population at withdrawal.

2.2 The population generator identity

For $x \in (0,1]^m$ put $V_x(z) = \prod_i x_i^{z_i}$. A type change, a death and a reproduction multiply V_x by x_j/x_i, by $1/x_i$ and by the offspring monomial divided by x_i, respectively. Summing conditional rates, the actual population generator obeys

$$\mathcal{G}_v V_x(z) = V_x(z) \sum_i z_i \frac{\Phi_v(x)_i}{x_i}. \tag{3}$$

This is an identity on the population state space. It remains valid under feedback because it is conditional on the current predictable action, and it makes no assertion that cells exposed to a common controller have independent histories; it does not factor founder probabilities. Conditional independence is used later only for a fixed baseline continuation, once the population at withdrawal is given.

2.3 First moments

Write L for the expected offspring matrix, $L_{ij} = \mathbb{E}[\#\{\text{daughters of type } j\} \mid \text{mother of type } i]$, so that $L = 2\Lambda$ for binary division with single-daughter marginal Λ. The first-moment matrix at constant action v is

$$A_v = Q_v + \operatorname{diag}(b)\,(L - I) - \operatorname{diag}(d), \tag{4}$$

acting on weight vectors from the left index, so that $W(z) = \sum_i z_i w_i$ has drift $\sum_i z_i (A_v w)_i$. Equivalently, A_v is the linearization of Φ_v at $x = \mathbf{1}$ up to sign: $D\Phi_v(\mathbf{1}) = -A_v$. It is a Metzler matrix. The one-daughter marginal determines A_v but not the extinction probabilities; sibling correlation lives in the joint law D and is retained throughout.

2.4 Endpoints that must stay distinct

Write

$$s_\pi(T) = \mathbb{P}_\pi(|Z_T| > 0), \qquad r_\pi(T) = \mathbb{P}_\pi(\text{non-extinction under the specified continuation after } T).$$

Then $s_\pi(T) \geq r_\pi(T)$, because extinction is absorbing and the continuation has no immigration. A finite-time survival *upper* bound therefore also bounds eventual survival; a lower bound on eventual survival also lower-bounds finite-time survival, provided its continuation assumptions hold. Our principal continuation withdraws the additional eraser, restoring baseline erasure while retaining the background demographic rates. Withdrawing the background killing as well is a different experiment, and Section 6.3 keeps the two separate.

2.5 Units

Reaction, division and death rates have units of inverse time. The integrated exposure $B = \int v\,dt$ is dimensionless when v is an additional first-order rate. Neither B nor v is a drug mass or a concentration. A nondimensionalization choosing a reference division rate b_0, setting $\hat{t} = b_0 t$ and dividing all rates by b_0 leaves $\int v\,dt$ unchanged, so exposure statements transfer between time units; changing the clock alone does not calibrate the relative biology. For the synthetic source of Section 6, $b_0 = .1$ and time 100 spans ten mean division waiting times.

3 The remaining-exposure theorem

Lemma 3.1 (Supersolution segment)**.** *Suppose $q \in (0,1]^m$, $\Phi_{\text{base}}(q) \le 0$, and let $h = 1 - q$. For $0 \le a \le 1$, $q_a = 1 - ah$ satisfies $\Phi_{\text{base}}(q_a) \le 0$. In the binary case, exactly*

$$\Phi_{\text{base}}(1 - ah) = a\Phi_{\text{base}}(q) - a(1-a)\, b \odot D(h,h). \tag{5}$$

If $Rq \le ch$ for $c \ge 0$, then $Rq_a \le c(1 - q_a)$.

Proof. Conservation of Q and normalization $F(1) = 1$ give $\Phi_{\text{base}}(1) = 0$. For binary offspring, expand $D(1-ah, 1-ah)$ and subtract $a\Phi_{\text{base}}(1-h)$; the linear chemical and death terms cancel as in (5), and $D(h,h) \ge 0$ because $h \ge 0$ and the daughter law has nonnegative coefficients. For general bounded offspring, each monomial of $F_i(1-ah)$ is a product $\prod_{\ell=1}^{k}(1 - ah_{j_\ell})$, whose second derivative in the scalar a is a sum of terms each containing two nonnegative factors h_j and otherwise nonnegative factors; hence each monomial, and so $F_i(1-ah)$, is convex *as a function of the scalar* a on $[0,1]$. Its graph lies below the chord, so $F_i(1-ah) \le (1-a)F_i(1) + aF_i(q)$. The remaining terms are affine in a, whence $\Phi_{\text{base}}(1-ah) \le a\Phi_{\text{base}}(q) \le 0$. Finally $R(1-ah) = aRq$ because $R\mathbf{1} = 0$. □

Remark 3.2. The convexity used here is directional: it is convexity along the prescribed segment $1 - a(1-q)$, not joint convexity of the generating function on the cube. Joint convexity is false in general — the monomial x_1x_2 has an indefinite Hessian — and nonnegative mixed partial derivatives do not supply it. The homotopy is affine in a; its later dependence on remaining exposure is exponential.

Theorem 3.3 (Exposure floor under feedback)**.** *Assume the model of Section 2. Let $q \in (0,1]^m$ and $c \ge 0$ satisfy*

$$\Phi_{\text{base}}(q) \le 0, \qquad Rq \le c(1-q). \tag{6}$$

Suppose the declared continuation has extinction vector $q^{\text{off}} \le q$. Then for every admitted predictable policy and every initial configuration z,

$$\mathbb{P}_z(|Z_T| > 0) \;\ge\; r_\pi(T) \;\ge\; 1 - \prod_i \left[1 - e^{-cB}(1 - q_i)\right]^{z_i}. \tag{7}$$

The same conclusion holds at an almost surely finite adaptive withdrawal time if the pathwise budget holds until that time and the same continuation condition holds conditional on its history.

Proof. For remaining exposure $\beta \ge 0$ set

$$a = e^{-c\beta}, \qquad q_\beta = 1 - a(1-q), \qquad L(\beta, z) = 1 - V_{q_\beta}(z).$$

All coordinates of q_β are positive, even at $\beta = 0$. Lemma 3.1 gives $\Phi_v(q_\beta) \le cv(1 - q_\beta)$ for every admitted $v \ge 0$, and $\partial_\beta q_\beta = c(1 - q_\beta)$. Differentiating the finite product in (3) yields

$$(\mathcal{G}_v - v\partial_\beta)L(\beta, z) = V_{q_\beta}(z) \sum_i \frac{z_i}{q_{\beta,i}} \left[cv(1 - q_{\beta,i}) - \Phi_v(q_\beta)_i\right] \;\ge\; 0. \tag{8}$$

This also holds at $z = 0$, where the sum is empty.

Let $\beta_t = B - C_t$, which is nonincreasing and nonnegative by (2). Stop at bounded population and proposal count and at time T. The finite jump compensation identity and the finite-variation chain rule make $L(\beta_t, Z_t)$ a submartingale after stopping. Non-explosion removes the stopping sequence, and $0 \le L \le 1$ supplies dominated convergence; therefore $\mathbb{E}L(\beta_T, Z_T) \ge L(B, z)$.

The function L is nonincreasing in β. Conditional on the history at T, the continuation survival probability is $1 - \prod_i (q_i^{\text{off}})^{Z_i(T)}$, so pointwise

$$L(\beta_T, Z_T) \le L(0, Z_T) \le 1 - \prod_i (q_i^{\text{off}})^{Z_i(T)} \le \mathbf{1}_{\{|Z_T|>0\}}.$$

Taking expectations proves (7). Independence is used only for the declared uncontrolled continuation, conditional on its starting population. For an almost surely finite withdrawal time τ, first stop at $\tau \wedge k$; boundedness and eventual equality to the stopped process at τ allow $k \to \infty$, after which the same terminal comparison applies. A rule that waits for extinction is not admissible under this extension unless it is almost surely finite. □

Remark 3.4. The hypothesis $\Phi_{\text{base}}(q) \le 0$ makes q an upper bound for the *minimal* extinction fixed point: $V_q(Z_t)$ is a bounded supermartingale for the uncontrolled source, extinction by time t forces $V_q = 1$, and letting $t \to \infty$ gives $q^{\text{off}} \le q$ whenever the continuation also satisfies $\Phi_{\text{off}}(q) \le 0$. An arbitrary numerical root of $\Phi = 0$ cannot be substituted without identifying it as that minimal solution.

Corollary 3.5 (Preparations, actuators and uncertainty)**.** *For a random initial configuration, average the right side of (7) over its actual law; for a preparation known only to lie in a set, minimize that right side over the set. Suppose instead that*

$$\Phi_u(x) = \Phi_{\text{base}}(x) + \sum_k u_k \Psi_k(x), \qquad \Psi_k(x) = R_k x + \kappa_k \odot (1 - x), \quad R_k \mathbf{1} = 0,$$

all admitted rates are valid, $u_k \ge 0$, and $\Psi_k(q) \le c_k(1 - q)$ with $c_k \ge 0$. Then replace cB in (7) by $\sum_k c_k B_k$ for pathwise budgets $\int u_k \le B_k$. For a single cost budget $\int \sum_k a_k u_k \le B$ with $c_k \le c_* a_k$, use $c_* B$. The statements are uniform over a fixed source parameter set if the same certificates and continuation inequalities hold throughout that set.

Proof. Condition on the preparation and apply Theorem 3.3. Linearity gives $\Psi_k(q_a) = a\Psi_k(q)$, so the segment bound of Lemma 3.1 holds with $\sum_k c_k u_k$ in place of cv. Use the remaining weighted exposure $\beta = \sum_k c_k B_k - \int \sum_k c_k u_k$ and $a = e^{-\beta}$ in (8). The cost-budget version uses $\sum_k c_k u_k \le c_* \sum_k a_k u_k$. The proof holds for each fixed source separately, giving uniformity without exchanging the policy and source quantifiers. □

A zero c_k is a property of this certificate; it does not prove that the actuator is biologically useless. A mixture of initial populations must not be replaced by its mean composition, since the right side of (7) is not affine in z. For application of a data-selected *upper* risk certificate, a joint confidence set with coverage $1 - \alpha$ and a policy bounded by δ throughout that set gives unconditional deployment risk at most $\delta + \alpha$, by conditioning on coverage and using independent deployment; independent marginal intervals do not automatically provide the required joint coverage.

Proposition 3.6 (Sharpness over the general class)**.** *The functional form of* (7) *is attained in a limiting admissible comparison source. This does not by itself establish sharpness of the constants in any particular molecular source.*

Proof. Take one immortal baseline type, no reproduction, and controlled death intensity cv. Every constant $q \in (0,1)$ has $\Phi_{\text{base}}(q) = 0$ and $\Psi(q) = c(1-q)$; after withdrawal the source is immortal and $q^{\text{off}} = 0$. Under a deterministic schedule using exposure B, independent survival of the n cells gives exactly $1-(1-e^{-cB})^n$. Letting $q \downarrow 0$ in Theorem 3.3 recovers equality. No zero denominator is used. □

Proposition 3.6 leaves open whether the form is attained inside a source with genuine inheritance. Section 7.4 shows that it is, in the following sense: in the molecular source with the added-death actuator the ratio of certified sufficient to necessary exposure is bounded by 1.45 as $n/\delta \to \infty$, uniformly in the number of memory sites.

4 Amplitude floors

Theorem 3.3 prices a course of treatment. It says nothing about a course of unbounded length and unbounded cost, because its right side tends to zero as $B \to \infty$. The second obstruction concerns exactly that regime: an actuator that is too weak cannot eradicate however long it is used.

Theorem 4.1 (Amplitude floor)**.** *Let $\mathcal{V} \subseteq [0, v_{\max}]$ be the admitted action set and suppose $q \in (0,1]^m$ satisfies*

$$\Phi_v(q) \le 0 \qquad \text{for every } v \in \mathcal{V}. \tag{9}$$

Then for every predictable policy with values in $\mathcal{V}$, with no budget restriction, every initial configuration z and every $t \ge 0$,

$$\mathbb{P}_z(|Z_t| > 0) \;\ge\; 1 - \prod_i q_i^{z_i}, \tag{10}$$

and the same bound holds for the probability of eventual non-extinction, and for eventual non-extinction after any continuation with $\Phi_{\text{off}}(q) \le 0$. Because $v \mapsto \Phi_v(q)$ is affine, (9) *for $\mathcal{V} = [0, v_{\max}]$ is equivalent to the two inequalities $\Phi_{\text{base}}(q) \le 0$ and $\Phi_{v_{\max}}(q) \le 0$.*

Proof. By (3) and (9), $\mathcal{G}_v V_q(z) = V_q(z) \sum_i z_i \Phi_v(q)_i / q_i \le 0$ for every admitted v and every z, including $z = 0$. Localizing as in Theorem 3.3 and using $0 \le V_q \le 1$, the process $V_q(Z_t)$ is a bounded supermartingale under any predictable policy, so $\mathbb{E}V_q(Z_t) \le V_q(z) = \prod_i q_i^{z_i}$. Extinction by time t implies $V_q(Z_t) = 1$, whence $\mathbb{P}_z(|Z_t| = 0) \le \mathbb{E}V_q(Z_t) \le \prod_i q_i^{z_i}$. Letting $t \to \infty$ and using that 0 is absorbing gives the statement for eventual extinction. For a continuation with $\Phi_{\text{off}}(q) \le 0$, apply the same supermartingale bound to the continuation started from Z_T and take expectations. Affineness of $v \mapsto \Phi_v(q) = \Phi_{\text{base}}(q) + vRq$ gives the last sentence. □

Remark 4.2. Theorem 4.1 is the degenerate case $c = 0$ of (6), but it is not a corollary of Theorem 3.3 as stated there: condition (6) with $c = 0$ requires $Rq \le 0$, which forces $\Phi_v(q) \le \Phi_{\text{base}}(q) \le 0$ for all $v \ge 0$ and is therefore a strictly stronger requirement than (9) on a bounded action range. The two certificates are genuinely different objects: (6) constrains the direction R relative to $1-q$ and yields a bound decaying in B; (9) constrains the field at the endpoint of the action range and yields a bound that does not decay at all.

Remark 4.3 (Cell-specific actions). Theorem 4.1 holds verbatim when each cell carries its own predictable action $v \in \mathcal{V}$, because (3) sums per-cell contributions and each is separately nonpositive. Theorem 3.3, by contrast, is a statement about one common budget and must not be silently reinterpreted as a sum of cell-specific exposures.

4.1 Sharpness of the amplitude threshold

Define the critical amplitude of a source as

$$v_c = \sup\{v \geq 0 : \ \exists\, q \in (0,1)^m \text{ with } \Phi_{\text{base}}(q) \leq 0 \text{ and } \Phi_v(q) \leq 0\}.$$

For $v_{\max} < v_c$, Theorem 4.1 gives a positive survival floor for every policy. In the other direction, suppose the first-moment matrix (4) at a constant admissible action v_* has spectral abscissa $\mathrm{s}(A_{v_*}) < 0$. Then, by Proposition 5.1 below, the constant action v_* drives the population to extinction almost surely, and the exposure required to reach risk δ from n founders is finite and of order $\log(n/\delta)$. In an irreducible finite-type source with nondegenerate reproduction, $q^{(v)} < \mathbf{1}$ if and only if $\mathrm{s}(A_v) > 0$ [1, 13], and the minimal extinction vector $q^{(v)}$ of the constant-v source is itself a candidate in the supremum defining v_c whenever it also satisfies $\Phi_{\text{base}}(q^{(v)}) \leq 0$. Consequently

$$v_c \ \leq \ \inf\{v : \mathrm{s}(A_v) < 0\},$$

and the two coincide whenever the minimal extinction vector at the larger amplitude is also a baseline supersolution. Section 7 computes $\inf\{v : \mathrm{s}(A_v) < 0\}$ exactly in the molecular source and exhibits amplitude certificates at values just below it — for instance at $e = .09$ when $e_c = .092750\ldots$ at two sites, and at $e = .35$ when $e_c = .373695\ldots$ at six — while the construction is infeasible above it. We do not claim the two quantities are equal, only that no gap is visible at the resolution computed.

4.2 A separate finite-time floor

Proposition 4.4 (Finite-time floor)**.** *For a binary source in which type changes preserve cell count and all controlled death rates are at most $d_{\max} > 0$,*

$$\mathbb{P}(|Z_T| > 0) \geq 1 - \left(1 - e^{-d_{\max} T}\right)^n, \qquad T \geq -\frac{1}{d_{\max}} \log\left[1 - (1-\delta)^{1/n}\right] \tag{11}$$

is necessary for finite-time risk at most δ from n founders.

Proof. Let $x(t) = 1 - e^{-d_{\max}(T-t)}$ and $f(t,z) = 1 - x(t)^{|z|}$, with $f(t,0) = 0$. For $t < T$ and $k = |z| > 0$, chemical changes contribute zero because they preserve $|z|$, and

$$(\partial_t + \mathcal{G}) f = x^{k-1}(1-x)\Big(k d_{\max} - \sum_i z_i d_i + x \sum_i z_i b_i\Big) \ \geq \ 0.$$

Localization and boundedness give the submartingale inequality. At T, $f(T,z) = \mathbf{1}_{\{|z|>0\}}$ read as a polynomial value. Taking expectations proves the first bound and rearrangement the second. □

The time and exposure restrictions are simultaneous necessary conditions; their probability bounds are not multiplied. The time floor does not apply to eventual survival after an arbitrarily long continuation. It requires non-death reproduction to leave at least one descendant, as binary division does, and any zero-offspring replacement event must be counted as an effective death.

5 Exposure and duration laws

5.1 The necessary side

For n founders of a single type with $\eta = 1 - q_i > 0$ and $c > 0$, risk at most $\delta \in (0,1)$ necessarily requires, by taking the nth root in (7),

$$B \ \geq \ B_{\text{nec}}(n,\delta) := \frac{1}{c}\left[\log \frac{\eta}{1 - (1-\delta)^{1/n}}\right]_+, \tag{12}$$

with $[x]_+ = \max(x, 0)$. For $0 < \delta \le 1/2$,

$$\frac{\delta}{n} \;\le\; 1-(1-\delta)^{1/n} \;\le\; \frac{-\log(1-\delta)}{n} \;\le\; \frac{2\delta}{n}. \tag{13}$$

The first inequality follows from concavity of $x^{1/n}$ at $x = 1$, the second from $1 - e^{-y} \le y$, the third by integrating $(1-x)^{-1} \le 2$ on $[0, \delta]$. Hence $B_{\mathrm{nec}}(n, \delta) \ge c^{-1}[\log(\eta n/(2\delta))]_+$.

5.2 The sufficient side

Proposition 5.1 (Weighted drift)**.** *Let v_* be an admissible constant action and suppose $w > 0$ satisfies $A_{v_*} w \le -\gamma w$ with $\gamma > 0$, where A_{v_*} is (4). Then under the constant action v_*,*

$$\mathbb{P}_z(|Z_t| > 0) \;\le\; \frac{\sum_i z_i w_i}{w_{\min}}\, e^{-\gamma t}. \tag{14}$$

For n identical founders of type i this gives risk at most δ once $t \ge \gamma^{-1}\log(Cn/\delta)$ with $C = w_i/w_{\min}$, at exposure

$$B_{\mathrm{suff}} = \frac{v_*}{\gamma}\,\log\frac{Cn}{\delta}, \tag{15}$$

provided $v_{\max} \ge v_*$ and the horizon is at least $\gamma^{-1}\log(Cn/\delta)$. Extinction is absorbing, so the bound persists after any continuation without immigration.

Proof. $W(z) = \sum_i z_i w_i$ has drift $\sum_i z_i (A_{v_*} w)_i \le -\gamma W(z)$. Localization with the first-moment domination of Section 2 and Gronwall give $\mathbb{E}W(Z_t) \le W(z)e^{-\gamma t}$. On $\{|Z_t| > 0\}$ we have $W(Z_t) \ge w_{\min}$, so Markov's inequality gives (14). Substituting $t = \gamma^{-1}\log(Cn/\delta)$ and $B = v_* t$ gives the rest. □

Since A_v is Metzler, $A_v w \le -\gamma w$ for some $w > 0$ implies $\mathrm{s}(A_v) \le -\gamma$, and conversely a positive witness exists with γ arbitrarily close to $-\mathrm{s}(A_v)$ when A_v is irreducible; the two-sided Collatz–Wielandt bounds

$$\min_i \frac{(Aw)_i}{w_i} \;\le\; \mathrm{s}(A) \;\le\; \max_i \frac{(Aw)_i}{w_i} \qquad (w > 0,\ A \text{ Metzler}) \tag{16}$$

turn any positive rational vector into a certified spectral bracket [3]. We use (16) throughout Section 7; a diagonal similarity by $\mathrm{diag}(w)$ preserves the spectrum and produces a Metzler matrix whose row sums lie between the two displayed quantities.

Corollary 5.2 (Matching logarithmic order)**.** *Fix a source with $\eta > 0$, $c > 0$ as in (6) and an admissible constant action with $\gamma > 0$ as in Proposition 5.1. Define $B^*(n, \delta)$ as the infimum of pathwise budget bounds among policies attaining risk at most δ, allowing any finite deterministic horizon; either the finite-time or the declared eventual endpoint may be used. With the amplitude fixed so as to admit that action,*

$$B^*(n, \delta) = \Theta\big(\log(n/\delta)\big) \qquad \text{as } n/\delta \to \infty, \quad 0 < \delta \le 1/2.$$

Proof. Equations (12)–(13) give $B^* \ge c^{-1}[\log(\eta n/(2\delta))]_+$ and (15) gives $B^* \le (v_*/\gamma)\log(Cn/\delta)$. All constants are fixed independently of n and δ. The upper bound supplies a finite horizon for each pair. It makes no feasibility assertion at a fixed short horizon, and it is vacuous when no admissible constant action has $\gamma > 0$ — precisely the situation Theorem 4.1 converts into a positive survival floor. □

The two constants c^{-1} and v_*/γ need not agree, and closing the gap between them is a source-specific question. Section 7.4 exhibits an actuator in the molecular source for which the ratio of the two is below 1.45, uniformly in the size of the inherited state.

6 A molecular inheritance realization

At N sites a cell has a activating, r repressive and $s = N - a - r$ unmarked sites. Following the reader–writer motif and the companion chronological inheritance construction [9, 11], the four chemical rates are

$$\begin{aligned} &(a,r) \to (a+1,r): \ s(.01 + a/N), && (a,r) \to (a,r+1): \ s(.01 + r/N), \\ &(a,r) \to (a-1,r): \ a(e + r/(2N)), && (a,r) \to (a,r-1): \ r(e + a/(2N)). \end{aligned} \tag{17}$$

Division removes the mother. The ordered daughters receive (i,j) and $(a-i, r-j)$ old marks with probability $2^{-(a+r)}\binom{a}{i}\binom{r}{j}$ and are refilled to N sites with unmarked material; there is no extra dilution term and no artificial recovery interval. Division is $b = .1$, and death is .01 if $a > r$ and .3 otherwise, so the states with $a > r$ are the *protected* ones. Intrinsic erasure is $e = .01 + v$, so the eraser actuator has $R = \partial_e Q_e$ and amplitude $v_{\max} = e_{\max} - .01$. The time unit is synthetic and these rates are modelling assumptions, not measured phenotype-switching rates.

For $N = 2$ order the states as (UU, UR, RR, AU, AR, AA). Their mark counts are, respectively, $(0,0)$, $(0,1)$, $(0,2)$, $(1,0)$, $(1,1)$ and $(2,0)$. A complete source specification is

$$\begin{aligned} (Q_e f)_0 &= (f_1 - f_0)/50 + (f_3 - f_0)/50, \\ (Q_e f)_1 &= .51(f_2 - f_1) + .01(f_4 - f_1) + e(f_0 - f_1), \\ (Q_e f)_2 &= 2e(f_1 - f_2), \\ (Q_e f)_3 &= .51(f_5 - f_3) + .01(f_4 - f_3) + e(f_0 - f_3), \\ (Q_e f)_4 &= (e + .25)(f_1 + f_3 - 2f_4), \\ (Q_e f)_5 &= 2e(f_3 - f_5), \end{aligned}$$

with the ordered daughter law

$$D(x,x) = \left(x_0^2,\ x_0 x_1,\ \frac{x_0 x_2 + x_1^2}{2},\ x_0 x_3,\ \frac{x_0 x_4 + x_1 x_3}{2},\ \frac{x_0 x_5 + x_3^2}{2} \right)^T.$$

The one-daughter marginal follows from this law; the mean operator is its linearization at $\mathbf{1}$, not an independent-daughter replacement. All exact computations below use the pair law.

6.1 Exact exposure certificates at two sites

Take

$$q = (.94, .982, .987, .27, .735, .235)^T, \qquad c = 67/73. \tag{18}$$

Exact rational calculation gives $\Phi_{.01}(q) \le 0$ with minimum slack $1/5000$. The derivative of Q_e satisfies

$$Rq = \left(0,\ q_0 - q_1,\ 2(q_1 - q_2),\ q_0 - q_3,\ q_1 + q_3 - 2q_4,\ 2(q_3 - q_5)\right)^T,$$

and the slacks in $c(1-q) - Rq$ are exactly

$$\left(\frac{201}{3650},\ \frac{534}{9125},\ \frac{1601}{73000},\ 0,\ \frac{33669}{73000},\ \frac{9229}{14600} \right)^T. \tag{19}$$

The certificate therefore holds for every nonnegative control increment for which the bounded-action construction is made. The baseline continuation satisfies $q^{\text{off}} \le q$ by the supermartingale argument

Table 1: Two-site source, AA founders, target risk .01, amplitude at least .29. The sufficient horizon is part of the guarantee. Rounded values display the exact logarithmic formulas (12) and (21); they are not rounded admission tests.

n	Necessary B	Sufficient B	Sufficient T
1	4.725705	19.282156	66.490193
10	7.229568	25.641677	88.419575
100	9.737862	32.001198	110.348957
1,000,000	19.772955	57.439281	198.066485

following Theorem 3.3, which also covers any off-source with $\Phi_{\mathrm{off}}(q) \le 0$. For n founders of type AA, Theorem 3.3 gives

$$r_\pi(T) \;\ge\; 1 - \left[1 - \frac{153}{200} e^{-67B/73}\right]^n. \tag{20}$$

The zero slack at AU makes $67/73$ the smallest admissible coefficient for this particular q and these six inequalities. It identifies a limiting certificate constraint, not a dominant biological trajectory and not an optimal intervention target.

Remark 6.1 (How much is lost by the displayed certificate). The bound (20) improves on the horizon-independent form $[1-(47/200)^n]e^{-67B/73}$ used in earlier drafts: writing $\alpha = e^{-cB}$, the function $1-(1-\alpha\eta)^n$ is concave in α and lies above its chord $\alpha[1-(1-\eta)^n]$, with equality at $n=1$, so the improvement is in the founder dependence and not a new one-founder constant. One may also ask how much is lost by fixing (18). Optimizing the certificate numerically and then verifying the result in exact rational arithmetic gives $\eta = .77849912\ldots$ and $c = .916533$ in place of $\eta = .765$ and $c = .917808\ldots$, raising $B_{\mathrm{nec}}(1, .01)$ from 4.725705 to 4.751365, an improvement of 0.54%. The vector is recorded in the data file `site_certificates.json`. The gap to the certified sufficient exposure of Section 6.2 is therefore not an artefact of the choice (18).

For the constant erasure $e = .3$, set $w_* = (1419, 1105, 1000, 8423, 4305, 10765)^T$. The exact residuals of $-(21/200)w_* - A_{.3}w_*$ are

$$(201/200,\ 37/40,\ 11/20,\ 47/8,\ 99/40,\ 301/40)^T > 0,$$

so Proposition 5.1 applies with $\gamma = 21/200$ and $C = 10.765$; equation (14) gives $10.765\, n e^{-.105t}$. With $v_{\max} \ge .29$, an early .29 increment for

$$\tau = \frac{200}{21} \log \frac{10.765n}{\delta}, \qquad B_{\mathrm{suff}} = \frac{58}{21} \log \frac{10.765n}{\delta} \tag{21}$$

achieves risk at most δ at every $T \ge \tau$ and after any subsequent continuation without immigration. This certificate is for $e = .3$ exactly, not for an unstated interval of uncertain erasure rates.

6.2 A certified finite-time pulse

For one AA founder use $e = .3$ on $[0, 40]$, then $e = .01$ on $(40, 100]$, with the same background killing throughout. The exposure is $B = .29 \times 40 = 11.6$. The backward calculation is $\mathcal{F}_{.3}^{40}(\mathcal{F}_{.01}^{60}(0))$, where $\mathcal{F}_e^t$ is the six-dimensional flow of $\dot u = \Phi_e(u)$. An outward-rounded dyadic Taylor enclosure proves

$$\mathbb{P}_{AA}(|Z_{100}| > 0) \;\le\; \frac{248922570493549955360831162 61}{25353012004564588029934064107 52} \;<\; 0.009818264214475 \;<\; .01. \tag{22}$$

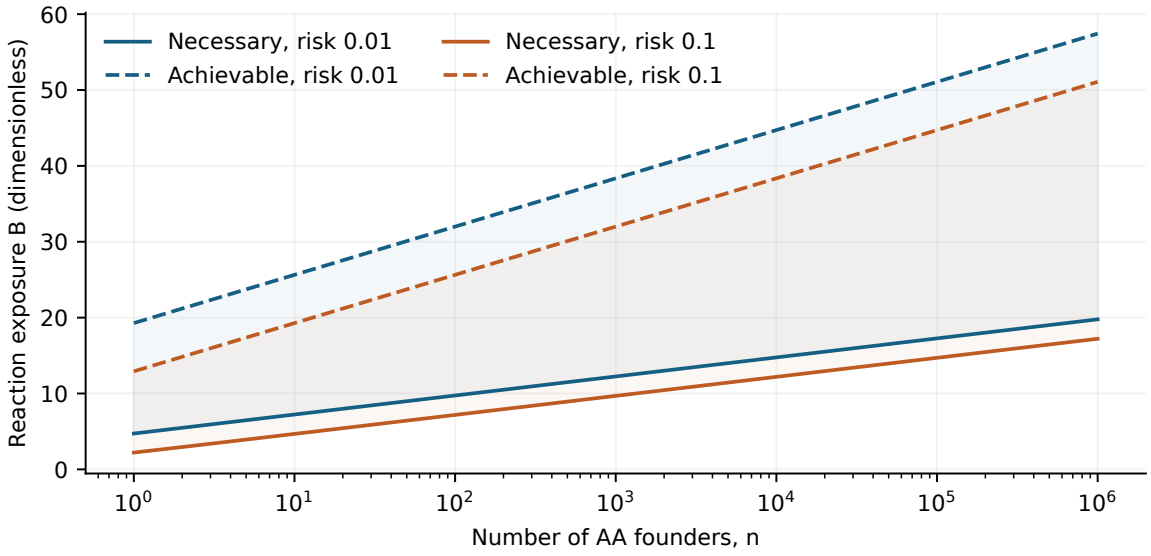

Figure 1: Proved exposure floor and achievable weighted bound for the two-site source, with free finite horizon. The gap is unclassified, not a numerically estimated optimal frontier. Exposure is reaction rate integrated over synthetic time.

Appendix A gives the error argument and arithmetic specification. The sufficient exposure at this particular horizon is therefore 11.6, against 19.282156… from (21) and a necessary 4.725704… from (12); the certified bracket at $n = 1$, $T = 100$, $v_{\max} = .29$, $\delta = .01$ is

$$4.725704\ldots \;\le\; B^*(100, .01) \;\le\; 11.6,$$

a ratio of about 2.45 rather than the 4.08 obtained from the mean certificate. Neither endpoint is asserted to be the exact optimum. For independent founders under a *deterministic* schedule, a single-founder upper risk ε gives population risk at most $1-(1-\varepsilon)^n \le n\varepsilon$; this use of independence is legitimate precisely because the schedule is predetermined. Equation (22) alone therefore does not achieve .01 for ten or a hundred founders.

6.3 Preparation and withdrawal are part of the theorem

Three experiments must be distinguished: eraser withdrawal under continuing background killing; withdrawal of both interventions into a specified off-source; and a later rechallenge or regrowth-threshold assay. The last requires its own stopped event and is not analyzed here. After eraser-only withdrawal our default off-source is $e = .01$ with the original death vector. If the off-source preserves Q, b, D and reduces death coordinatewise then $\Phi_{\text{off}}(q) \le \Phi_{\text{base}}(q) \le 0$ and the exposure floor persists. For the explicit full-withdrawal scenario with $b = .1$, $d = .01$ in every state, the total count is an ordinary birth–death process with extinction vector $q^{\text{off}} = .1\,\mathbf{1}$, the smaller root of $.01(1-x) + .1(x^2 - x) = 0$. This is a stipulated full-withdrawal model, not a fitted biological law.

Erasure is not universally beneficial at a fixed short horizon. From an RR founder the analytic finite-time extinction solution satisfies

$$u^{3}_{RR}(t;0) - u^{01}_{RR}(t;0) = -\frac{49619}{600000000}\,t^4 + O(t^5), \tag{23}$$

so stronger erasure *lowers* extinction probability for all sufficiently small $t > 0$. This is a source-level failure of monotonicity: erasure also removes repressive marks and permits access to the protected

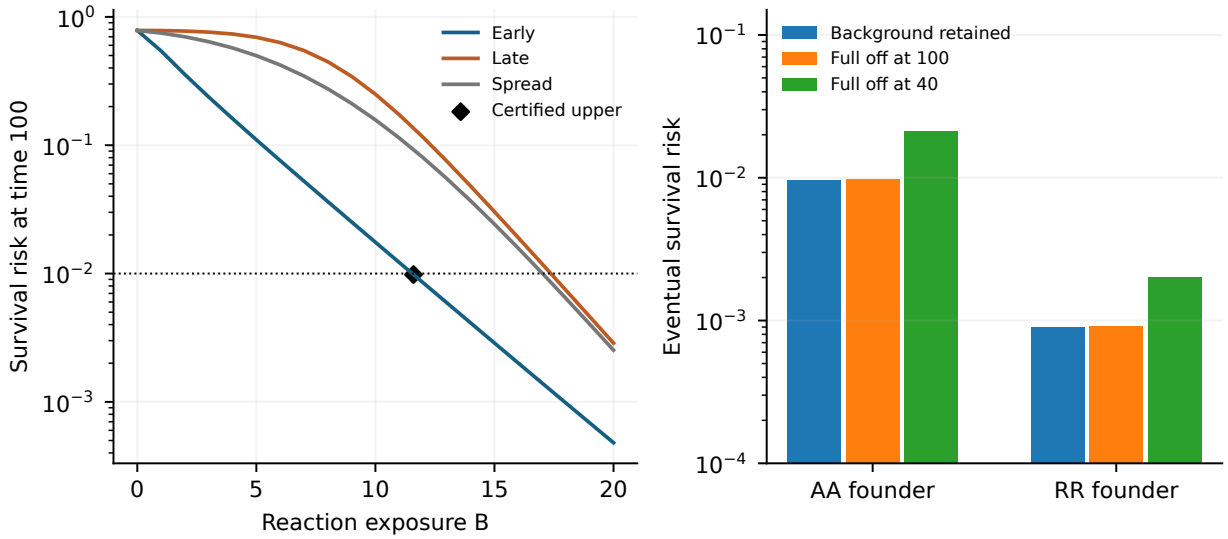

Figure 2: Left: AA survival by time 100 for equal-exposure early, late and spread eraser schedules (lines: numerical diagnostics). The marked upper bound at $B = 11.6$ is rigorous. Right: numerical eventual-survival comparisons after eraser withdrawal at 40, keeping background killing until 100 (then baseline or full withdrawal), or removing all killing immediately at 40. Every plotted continuation is stated; no curve is a clinical relapse estimate.

active states, and the source assigns lower death to $a > r$. It is a rigorous counterexample to universal monotonicity, not evidence that an epigenetic drug rescues tumours, and it does not imply a large experimental effect: at $t = 1$ the numerical difference is about 4.1×10^{-5}. The exact calculation is in Appendix B; finite-time numerical differences are labelled separately. Preparation, budget, duration and endpoint all matter, and an early/late ordering seen for AA need not persist for RR.

6.4 Slower protected division

To challenge the demographic assumptions without rescaling the clock, replace the division rate in states $a > r$ by a common unknown fixed $b_P \in [.08, .1]$, keeping all death and chemical rates fixed. An independently checked rational supersolution and positive weight give, uniformly on that interval,

$$q_{AA} = .34796677, \qquad c = 0.9163811468\ldots, \qquad \gamma = .1004, \qquad C = 10.764704.$$

The data file `source_certificates.json` records the exact value of c, all coordinates and the endpoint residuals. Each inequality is affine in b_P, so exact endpoint checks prove the whole interval. This is a scenario family, not a confidence interval inferred from experiments, and a 20% reduction in protected division is a modest perturbation; it does not demonstrate robustness to deep dormancy or to orders-of-magnitude slower persister division.

7 How the obstructions scale with memory size

Everything in Section 6 concerns two sites. The construction (17) is defined for every N, with $\binom{N+2}{2}$ states, and the question is which of the two obstructions grows with N. This section answers that

Table 2: Exact certificates for the N-site source (17). The critical erasure $e_c(N)$ is bracketed by the M-matrix test (24) at denominator 10^6; $v_c = e_c - .01$ is the corresponding actuator amplitude. The last two columns are a positive rational weight certificate $A_{1/2}w \le -\gamma w$ at the common admissible erasure $e = 1/2$, with $C = w_{(N,0)}/w_{\min}$ as in Proposition 5.1. All rows are exact finite computations; γ is exact at six decimals and C is rounded up.

N	states	$e_c(N) \in$	$v_c(N)$	γ at $e = 1/2$	C at $e = 1/2$
2	6	$(.092750, .092751)$	.0827	.166376	3.7420
3	10	$(.188128, .188129)$	.1781	.129493	6.4630
4	15	$(.263685, .263686)$	.2537	.095088	9.3007
5	21	$(.324157, .324158)$	.3142	.067868	11.9061
6	28	$(.373695, .373696)$	.3637	.046787	14.2792
7	36	$(.415181, .415182)$	.4052	.030242	16.4792
8	45	$(.450572, .450573)$	.4406	.017016	18.5583

for $2 \le N \le 8$ with exact rational certificates: the exposure floor is almost flat in N, while the *amplitude* required by the eraser grows by a factor of nearly five. In particular the amplitude used in the two-site worked example is certified insufficient, at any budget and any duration, at five to eight sites.

7.1 An exact test for the sign of the mean growth rate

Let A_e be the first-moment matrix (4) of the N-site source at constant intrinsic erasure e, reconstructed from (17), the binomial allocation law and the death rule. It is Metzler and, for $e > 0$, irreducible, because basal writing and intrinsic erasure connect every state to every other. Then $M = -A_e$ is a Z-matrix, and the standard equivalence for Z-matrices [3, Ch. 6] gives

$$\begin{aligned} \mathrm{s}(A_e) < 0 \quad &\iff \quad M \text{ is a nonsingular } M\text{-matrix} \\ &\iff \quad \text{all leading principal minors of } M \text{ are positive.} \end{aligned} \tag{24}$$

The right-hand condition is decided exactly by fraction-free elimination over $\mathbb{Q}$, and the rates in (17) are rational whenever e is. Bisecting on e with denominator 10^6 therefore yields a rigorous bracket for the critical erasure

$$e_c(N) := \inf\{e > 0 : \ \mathrm{s}(A_e) < 0\},$$

which is well defined because $e \mapsto \mathrm{s}(A_e)$ changes sign exactly once on the tested range. The brackets are the first two numeric columns of Table 2; each was obtained by 21 exact sign evaluations and each is verified independently by a positive rational Collatz–Wielandt witness at both endpoints, recorded in the data file `site_certificates.json`.

The two-site value reproduces, to the displayed resolution, the unique positive root of the determinant polynomial

$$8{,}000{,}000{,}000\, e^5 + 7{,}688{,}000{,}000\, e^4 + 2{,}886{,}960{,}000\, e^3 + 439{,}939{,}800\, e^2 - 16{,}495{,}365\, e - 5{,}182{,}134, \tag{25}$$

which equals $5 \cdot 10^9 \det A_e$ at $N = 2$; its unique positive real root is $.09275086931\ldots$, and the remaining four roots have negative real part.

Table 3: Certified amplitude floors: exact rational q with $\Phi_v(q) \le 0$ for all $v \in [0, v_{\max}]$, so that no policy of that amplitude eradicates at any budget or horizon. $\eta_{(N,0)} = 1 - q_{(N,0)}$ bounds survival from one fully activating founder, and $1 - \max_i q_i$ bounds it from the least favourable founder type. The vectors were found by iterating the componentwise maximum of the two fields to its minimal fixed point, retracted by the factor 99/100 along the segment of Lemma 3.1 and then verified exactly; the retraction costs 1% of η. The last row shows that the construction remains feasible at $v_{\max} = .34$, close to the certified threshold $v_c(6) = .363695\ldots$ of Table 2.

N	$v_{\max}$	$\eta_{(N,0)}$	$1 - \max_i q_i$
2	.02	.67850219	.01263148
2	.05	.42726916	.00771882
2	.08	.03785729	.00067026
5	.29	.17008641	.00185216
6	.29	.45962721	.00434343
7	.29	.61348974	.00504729
8	.29	.70130311	.00503641
6	.34	.14866289	.00133682

7.2 No policy eradicates below the threshold

Proposition 7.1 (Certified amplitude floors)**.** *Consider the source* (17) *with baseline erasure* $e = .01$ *and eraser amplitude* $v_{\max}$*, and let* $\eta_{(N,0)}$ *be the value in Table 3. For every predictable policy with values in* $[0, v_{\max}]$*, including cell-specific policies, with no budget restriction and no horizon restriction, survival from* n *founders of the fully activating type* $(N, 0)$ *satisfies*

$$\mathbb{P}(\textit{non-extinction}) \;\ge\; 1 - \big(1 - \eta_{(N,0)}\big)^n,$$

and survival from a single founder of any type is at least $1 - \max_i q_i$*. In particular, at the amplitude* $v_{\max} = .29$ *used in Section 6.1, no policy whatsoever eradicates the* N*-site source for* $5 \le N \le 8$*.*

Proof. The explicit rational vectors q in the data file `site_certificates.json` satisfy $\Phi_{\text{base}}(q) \le 0$ and $\Phi_{v_{\max}}(q) \le 0$ with strictly negative slacks in every coordinate, verified in exact arithmetic; the largest slack in each case is reported in the data file. Theorem 4.1 and Remark 4.3 then give the two displayed bounds. □

Two points deserve emphasis. First, the amplitude floors of Table 3 are not statements about a particular constant action: they hold against every predictable feedback policy, including ones that exploit the non-monotonicity of (23) by modulating the eraser, and they hold at infinite budget. A mean-growth diagnostic at a single constant action — for instance the exact bracket $\mathrm{s}(A_{\cdot 3}) \in [239/10000, 3/125]$ at $N = 6$, or $\mathrm{s}(A_{\cdot 3}) \in [-7/500, -139/10000]$ at $N = 4$, both certified by positive rational Collatz–Wielandt witnesses — shows only that *that* action fails to contract, and leaves the policy question open. Theorem 4.1 closes it.

Second, the failure at $5 \le N \le 8$ is a property of the source and not of a loose certificate. The published four-site weight of the earlier two-site study achieved $\gamma = .0132$ at $e = .3$, which is at least $.0132/.014 > 94\%$ of the true mean decay rate at that action; almost none of the deterioration between two and four sites is recoverable by searching harder for a linear weight.

Remark 7.2 (Beyond the certified range). Floating eigenvalue computations of $\mathrm{s}(A_e)$, exact in the model but not rounded rigorously, continue the first column of Table 2 as $e_c \approx .508$ at $N = 10$, .621

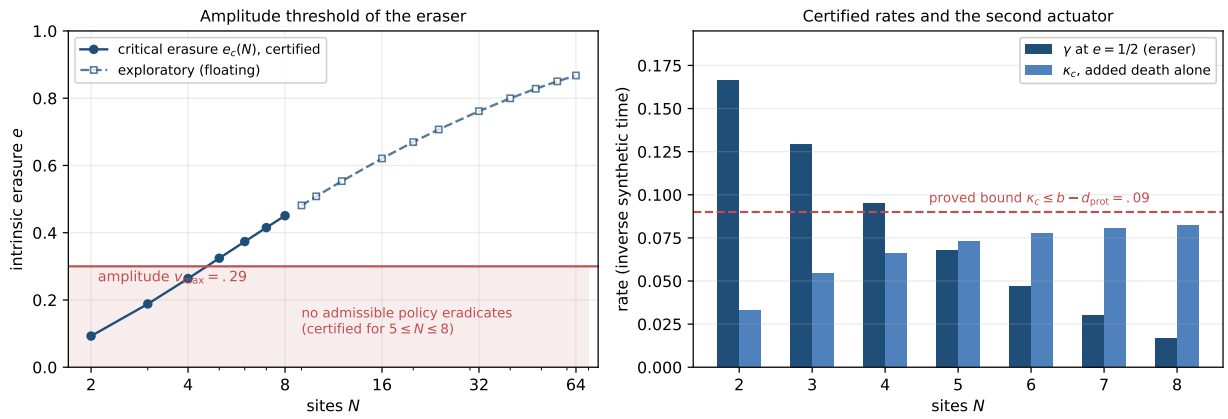

Figure 3: Left: the eraser's critical intrinsic erasure against memory size. Filled markers are the exact brackets of Table 2; open markers are exploratory floating eigenvalue computations for $9 \le N \le 64$ and are diagnostics, not certificates. The shaded band is the amplitude used in the two-site worked example; for $5 \le N \le 8$ the amplitude floors of Table 3 certify that no policy in that band eradicates. Right: certified mean contraction rate γ of the eraser at the common erasure $e = 1/2$, and the certified threshold κ_c of the added-death actuator acting alone, which Proposition 7.3 bounds by $b - d_{\text{prot}} = .09$ for every N.

at $N = 16$, .762 at $N = 32$ and .868 at $N = 64$. The increments per unit of $\log N$ fall from about .27 to about .13 over that range, so the sequence is increasing with decelerating increments and we do not determine whether $e_c(N)$ remains bounded. What is proved here is the monotone increase and the factor 4.86 between $N = 2$ and $N = 8$, together with the amplitude floors of Table 3; no assertion is made uniformly in N for the eraser, and six or eight abstract sites are not claimed to be a realistic chromatin domain.

7.3 The exposure floor barely moves

Above the amplitude threshold the picture is different. Table 4 gives, for each N, an exact baseline supersolution and the smallest admissible c for the eraser; the resulting necessary exposure at $n = 1$, $\delta = .01$ increases by only 22% between two and eight sites, against the factor 4.86 in amplitude. Larger inherited state makes the actuator harder to *power*, not appreciably more expensive to *run*, in this source.

The constructive side does deteriorate at a fixed amplitude: at the common admissible erasure $e = 1/2$, Table 2 gives γ falling from .166 to .017 and C rising from 3.7 to 18.6, so (15) yields $B_{\text{suff}}(1, .01)$ growing from 17.4 at $N = 2$ to 216.8 at $N = 8$. Raising the amplitude helps: a floating optimization of $v_*/\gamma(v_*)$ over constant actions gives best free-horizon exposures of about 17.1 at $N = 2$ (at $e = .43$), 27.8 at $N = 4$ (at $e = .74$) and 34.4 at $N = 6$ (at $e = .91$). These are diagnostics; the certified rows are Table 2. Either way the eraser's sufficient exposure grows with N while its necessary exposure does not, so the constant gap widens and the question of the true optimum becomes harder, not easier, with memory size.

Table 4: Exact baseline exposure certificates for the N-site source: $\Phi_{\text{base}}(q) \le 0$ and $Rq \le c(1-q)$, with $\eta_{(N,0)} = 1 - q_{(N,0)}$ and B_{nec} evaluated from (12) at $n = 1$, $\delta = .01$. The $N = 2$ row is the optimized certificate of Remark 6.1, which supersedes (18) by 0.54%. All vectors are in the data file `site_certificates.json`.

N	$\eta_{(N,0)}$	c	$B_{\text{nec}}(1, .01)$
2	.77849912	.916533	4.751365
3	.85886565	.884640	5.033717
4	.88009489	.857113	5.223867
5	.88886370	.832987	5.387069
6	.89308829	.811604	5.534843
7	.89533756	.792438	5.671883
8	.89661409	.775121	5.800437

7.4 A second actuator removes the amplitude obstruction

Corollary 3.5 already admits several additive actuators. The natural second one here adds death on the protected states,

$$\Psi(x) = \kappa \odot (1 - x), \qquad \kappa_i = \mathbf{1}_{\{a_i > r_i\}}, \tag{26}$$

with control $u \in [0, u_{\max}]$, so that the death rate in protected states becomes $d_{\text{prot}} + u$. This is a defined mathematical actuator, not a claim that a known drug has that selectivity. Because $\kappa \le \mathbf{1}$ we may take $c_\kappa = 1$ in Corollary 3.5, so the necessary exposure for this actuator is exactly $\log(\eta/(1-(1-\delta)^{1/n}))$ with η from Table 4.

Proposition 7.3 (Dimension-free sufficiency)**.** *Let a finite-type source have binary division at rate b, death rates d_i, count-preserving type changes, and an added-death actuator with nonnegative indicator vector κ as in (26), held at a constant value $u = \kappa_*$. If*

$$\gamma := \min_i (d_i + \kappa_* \kappa_i - b) > 0,$$

then $\mathbb{E}|Z_t| \le |z| e^{-\gamma t}$, so $\mathbb{P}_z(|Z_t| > 0) \le |z| e^{-\gamma t}$ and extinction is almost sure. In the source (17), where $b = .1$, $d_{\text{prot}} = .01$ and $d_{\text{unprot}} = .3$, this holds for every N as soon as $\kappa_* > b - d_{\text{prot}} = .09$, with $\gamma = \min(\kappa_* - .09,\ .2)$ and $C = 1$; hence

$$B_{\text{suff}} = \frac{\kappa_*}{\min(\kappa_* - .09,\ .2)} \log \frac{n}{\delta}, \tag{27}$$

minimized at $\kappa_* = .29$, where it equals $1.45 \log(n/\delta)$ for every number of sites.

Proof. Type changes preserve $|z|$ and binary division replaces one cell by two, so $|Z_t|$ increases by one at rate $b|Z_t|$ and decreases by one at rate $\sum_i Z_i(t)(d_i + \kappa_* \kappa_i)$. Hence $\frac{d}{dt}\mathbb{E}|Z_t| = \mathbb{E}\sum_i Z_i(t)(b - d_i - \kappa_* \kappa_i) \le -\gamma \mathbb{E}|Z_t|$, after the localization of Section 2, and Gronwall gives the bound. Since $|Z_t|$ is integer valued, $\mathbb{P}(|Z_t| > 0) \le \mathbb{E}|Z_t|$. This is Proposition 5.1 with $w = \mathbf{1}$, so $C = 1$; substituting $t = \gamma^{-1} \log(n/\delta)$ and $B = \kappa_* t$ gives (27). In (17) the protected states have $b - d = .09$ and the unprotected ones $b - d = -.2$, giving the stated γ. Minimizing $\kappa_* / \min(\kappa_* - .09, .2)$ over $\kappa_* > .09$ gives the value $.29/.2 = 1.45$ at $\kappa_* = .29$. □

Corollary 7.4 (Near-sharp exposure constant)**.** *For the actuator* (26) *in the source* (17) *with amplitude* $u_{\max} \geq .29$*, at every* $N \leq 8$*,*

$$\log \frac{\eta_{(N,0)}}{1-(1-\delta)^{1/n}} \;\leq\; B^*(n,\delta) \;\leq\; 1.45 \log \frac{n}{\delta},$$

so $\limsup_{n/\delta\to\infty} B^*(n,\delta)/\log(n/\delta) \leq 1.45$ *while the liminf is at least* 1*. At* $n = 1$*,* $\delta = .01$ *the certified bracket is* $4.354788 \leq B^* \leq 6.677497$*, a ratio of* 1.533*.*

Proof. Combine Corollary 3.5 with $c_\kappa = 1$ and the values of η in Table 4 for the lower bound, and Proposition 7.3 for the upper. For the liminf use (13), which gives $B^* \geq \log(n/\delta) + \log(\eta/2)$. □

This is the sharpness statement that Proposition 3.6 could not supply inside a source with genuine inheritance: the exposure floor (7) is attained to within a factor 1.45 by an explicit policy in the molecular source, uniformly in the size of the inherited state, and the leading order $\log(n/\delta)$ is exact.

Finally, a small amount of the second actuator repairs the eraser's failure rather than replacing it. With the eraser held at $e = .3$ — the amplitude that Proposition 7.1 certifies cannot eradicate on its own for $5 \leq N \leq 8$ — the added-death threshold is bracketed exactly at denominator 10^6 by

$$\kappa_c \in (.008894, .008895),\ (.024937, .024938),\ (.036344, .036345),\ (.044765, .044766)$$

for $N = 5, 6, 7, 8$ respectively. When the eraser is absent the corresponding thresholds are

$$(.033181,\ .054493,\ .066139,\ .073090,\ .077472,\ .080374,\ .082373)$$

for $N = 2, \ldots, 8$ — an increasing sequence that Proposition 7.3 caps at .09 for every N. At six sites an added death of .025 on protected states therefore restores mean contraction where an eraser amplitude of .29 provably cannot. Additive certificate costs do not by themselves prove or disprove biological synergy; synergy must be defined against a stated cost, dose pair, endpoint and admissible policy class, and non-additive target interactions require their own generator bounds.

8 Interpretation and an experimental test

8.1 Parameter provenance and observation

Every numerical rate in Section 6 is assumed. Basal writing, feedback, recruited erasure and division allocation specify an inherited-state mechanism; they are not measured phenotype-switching rates. The phenotype-to-death map is an explicit functional assumption, and the larger-site sources of Section 7 are further mathematical realizations, not more accurately calibrated loci. A physical larger domain might use a different protected-state threshold, enzyme limitation, cooperative recruitment or site dependence; those would be different sources and should be named as such.

Experiments would need separate division and death event records, calibrated target engagement, a state observation model and *joint* daughter observations. Net growth $b - d$ alone is insufficient: the single-type sources $(b, d) = (.1, .05)$ and $(.2, .15)$ share net growth .05 but have eventual extinction probabilities 1/2 and 3/4. A binary reporter has emission rank at most two and cannot satisfy full-column-rank identification of six arbitrary states; this blocks that particular identification argument, not the possibility of identifying a constrained low-dimensional rate family or a certificate-relevant functional. No numerical range from another cell line or treatment model is transplanted into this source.

8.2 Conditional concentration and administered-amount bounds

Suppose the target-engaging concentration $C(t) \geq 0$ has a calibrated saturating effect

$$v(C) = v_{\max}\frac{C}{K+C}, \qquad K, v_{\max} > 0.$$

Here C and K carry concentration units, $v_{\max}$ inverse time, and the reaction exposure $B_{\text{actual}} = \int_0^T v(C(t))\,dt$ is dimensionless. For finite concentration AUC $A = \int_0^T C(t)\,dt$, concavity and Jensen give

$$B_{\text{actual}} \leq \frac{v_{\max}TA}{KT+A}.$$

If (12) requires $B_{\text{actual}} \geq B_{\text{req}} > 0$, then

$$v_{\max}T > B_{\text{req}}, \qquad A \geq \frac{KTB_{\text{req}}}{v_{\max}T - B_{\text{req}}}. \tag{28}$$

The strict inequality holds because finite A makes $A/(KT+A) < 1$; multiplication by positive denominators proves the second. A concentration cap adds $B_{\text{req}} \leq Tv(C_{\max})$.

For $\dot{C} = I/V_d - kC$ with $C(0) = 0$, nonnegative infusion I, elimination $k > 0$ and distribution volume $V_d > 0$, integration gives $M = V_d(C(T) + kA) \geq kV_dA$, where $M = \int I$ is the administered amount, so

$$M \;\geq\; \frac{kV_dKTB_{\text{req}}}{v_{\max}T - B_{\text{req}}}. \tag{29}$$

A nonzero initial concentration instead gives $M = V_d(C(T) - C(0) + kA)$ and subtracts at most $V_dC(0)$ from the lower bound, truncated at zero. These are necessary conditional bounds, not sufficient regimens; plasma concentration may not equal the target-engaging concentration; and for a Hill exponent above one global concavity fails, so one must use a valid concave majorant on the admitted range rather than reusing this Jensen argument. The inequality prices the eraser only, not the background killing retained in the example.

The useful content is qualitative and available before calibration: required engagement close to $v_{\max}T$ forces the necessary AUC to increase sharply. Section 7 sharpens the point. The amplitude floor is a constraint on $v_{\max}$ alone, and no amount of AUC substitutes for it: if $v_{\max} < v_c$, (28) is not merely hard to satisfy, it is irrelevant, because Theorem 4.1 rules out eradication at any A and any T.

8.3 A falsifiable low-density comparison

A prospective microcolony design can compare known founder counts $1, 10, 100$, with AA-rich and RR-rich preparations, equal reaction-exposure early, late and spread schedules, and separately declared eraser-only and complete-withdrawal arms. Division and death event tracking should accompany the molecular readout, with all founders retained in denominators. A fixed-time viable-cell census and outgrowth over a declared observation window are different events; undetected cells cannot be equated with extinct lineages without a validated detection model.

The prediction is a one-sided requirement on exposure, not equality with a fitted boundary. With a jointly valid parameter and preparation uncertainty set, the lower bound is minimized and any policy upper bound maximized over the set before comparing with data. For a simple fixed protocol and independent replicates, a binomial confidence interval for the correctly observed event can test that propagated prediction, with multiple-look error controlled in advance. Survival significantly below a valid lower bound would challenge the adopted source, actuator law or observation mapping; a data point above the lower bound, or away from the sufficient upper curve, does not falsify it.

8.4 Density, space and what the model is not

The branching model is naturally a low-density persistence or microcolony model. It is not automatically a model of a macroscopic tumour, and the large-n statements are mathematical scaling, not a claim that a million-cell experiment remains density independent. Some extensions reuse the barrier directly: for a population-dependent source it suffices to prove the analogue of (8) uniformly over the admitted population states and environments, and a constructive upper bound survives if the weighted drift remains uniformly negative. Reducing division is not automatically favourable here, because division also redistributes marks and the sign of $D_i(x, x) - x_i$ need not be uniform; density-dependent birth suppression therefore cannot be inserted by an unsupported monotonicity argument. Spatial migration or hidden carriers require explicit additional state variables and generator terms, and persistent immigration destroys the absorbing-zero setup entirely.

9 Discussion

The remaining-exposure construction establishes a population-dependent floor against arbitrary admitted feedback. The product in that floor is a function of the whole configuration, not an independence approximation during treatment, and its founder dependence survives the coupling that a common controller introduces. The amplitude construction is a different certificate with a different conclusion: below a critical actuator strength the survival floor does not decay at all, so no budget and no schedule can help. Separating the two is what makes the memory-size question well posed, and what turns a diagnostic about one constant action into a statement about every policy.

Three concrete conclusions follow for the molecular source. First, the obstruction that scales badly with the size of the inherited state is *amplitude*, not exposure: between two and eight sites the certified critical erasure rises by a factor 4.86 while the necessary exposure rises by 22%. Second, the amplitude used in the two-site worked example provably fails at five to eight sites, against every predictable feedback policy, at infinite budget and infinite time. Third, the failure is a property of that actuator: added death on protected states has a threshold bounded by $b - d_{\mathrm{prot}}$ for every N by an elementary total-count argument, delivers a dimension-free sufficient exposure $1.45 \log(n/\delta)$, and attains the exposure floor to within a factor 1.533 at $n = 1$, $\delta = .01$. That last point also supplies the sharpness that Proposition 3.6 could only give in a degenerate comparison source.

What remains open is quantitative and specific. For the eraser at two sites the certified bracket at $T = 100$ is $4.7257 \le B^* \le 11.6$, and Remark 6.1 shows the lower endpoint is within 0.54% of what certificate optimization alone can deliver, so the residual factor of about 2.45 is not a defect of the chosen supersolution. Closing it requires either a richer population barrier or a better policy, and the natural next experiment is to certify a small phase family selected by backward sensitivity of the generating-function flow and compare it with a population-budget barrier that does not factor through a single vector q. Whether $e_c(N)$ remains bounded as $N \to \infty$ is open; the exploratory evidence of Remark 7.2 is consistent with either answer, and a proof in either direction would say something structural about how fast inherited memory can be defeated by a single reaction-rate actuator.

Uniform constants in site count for the eraser, density-dependent populations, empirical target engagement, a global solution of the associated dynamic programming equation, and full stochastic formalization all remain outside this paper. They are explicit future scopes and not hidden assumptions: the formalization boundary is stated in Appendix D, the evidence type of every displayed quantity is labelled, and the exact vectors behind every certificate are released with the source package.

A Validated generating-function calculation

This appendix specifies the proof behind (22) independently of any floating ODE solver. For constant erasure write $\Phi(z) = Jz + d + .1D(z,z)$ with $J = Q - \mathrm{diag}(d + .1)$. Taylor coefficients at any initial point obey

$$c_0 = z, \qquad c_{k+1} = \frac{1}{k+1}\left[Jc_k + \mathbf{1}_{\{k=0\}}d + .1\sum_{j=0}^{k} D(c_j, c_{k-j})\right]. \tag{30}$$

The unit cube is invariant: at $z_i = 0$ every inward term is nonnegative, and at $z_i = 1$ the switching, death and offspring terms sum to a nonpositive value. On the cube the Jacobian is Metzler, with row sums $-d_i + .1\big(\sum_j \partial_j D_i(z,z) - 1\big) \le .1 - d_i \le .09$. The upper right derivative of the sup norm of a difference of solutions is therefore at most .09 times that norm, by the mean-value Jacobian and the maximal-coordinate argument, so the flow is $e^{.09t}$-Lipschitz in that norm. This is a one-sided logarithmic-norm bound [19], not a bound on the norm of the Jacobian itself.

All interval endpoints in `scripts/validated_pgf.py` are integers divided by 2^{110}. Addition and subtraction are exact at that scale. Multiplication takes the smallest and largest of the four integer products and rounds outwards after division by 2^{110}; division by a positive integer also rounds outwards; rational inputs are enclosed by floor and ceiling. Induction on (30) proves inclusion of each exact coefficient [14].

Use steps $h = 1/20$ and a degree-11 Taylor polynomial. First evaluate the recurrence to c_{12} with $c_0 = [0,1]^6$ and let M be the largest absolute endpoint; at every point of the exact trajectory $\|z^{(12)}\|_\infty/12! \le M$, so Taylor's integral remainder on a step is bounded by Mh^{12}. The bounds used are approximately $4.828 \cdot 10^{-6}$ for $e = .01$ and $1.806 \cdot 10^{-4}$ for $e = .3$; the exact fractions are in the data file `validated_policy.json`.

For each step, compute the polynomial at the dyadic centre using outward interval operations, add the remainder, and choose a dyadic midpoint projected onto the unit cube. The maximum distance from that centre to the interval, plus the remainder, bounds the local error; this local enclosure is defined once and its radius enters the error recurrence once. If E bounds the incoming sup-norm error, propagate it as

$$E_{\mathrm{new}} \;\le\; \frac{E}{1 - .09h} + E_{\mathrm{local}},$$

rounded upward, which is valid because $e^{.09h} \le (1 - .09h)^{-1}$ for $.09h < 1$. Projection does not invalidate the calculation: the local distance is computed after projection and the exact flow remains in the cube. Starting at zero, 1200 low-erasure steps followed by 800 high-erasure steps give total error

$$E = \frac{3241002892060825847}{1298074214633706907132624082305024} \;<\; 2.497 \cdot 10^{-15}.$$

The enclosure accounts separately for coefficient rounding, polynomial evaluation, local truncation, error propagation and phase composition. It is an exact-arithmetic computation together with the preceding conventional enclosure proof [15], not a formalization of the numerical algorithm in a proof assistant. No unproved floating tolerance enters the certificate.

B The preparation counterexample

For $u(0) = 0$ the first Taylor coefficients of the backward flow are $c_1 = d$, $c_2 = J_e d/2$, $c_3 = (J_e c_2 + .1D(d,d))/3$ and $c_4 = (J_e c_3 + .2D(d,c_2))/4$. Substituting in the RR coordinate of the

two-site source gives

$$(c_4)_{RR} = -\frac{16087}{8000000} - \frac{29}{480000}\,e - \frac{29}{40000}\,e^2.$$

The coefficients through order three agree between $e = .01$ and $e = .3$, and their fourth-order difference is the coefficient displayed in (23):

$$(c_4)_{RR}\big|_{e=.3} - (c_4)_{RR}\big|_{e=.01} = -\frac{49619}{600000000}.$$

Analyticity of a polynomial ODE near zero then proves the strict sign on some positive interval, which is the content of (23). Numerically, at $t = 1$ the two extinction probabilities are .2493000462 and .2492590475, a difference of about $4.1 \cdot 10^{-5}$; these displayed finite-time values are diagnostics and not part of the analytic proof. The example shows that stronger erasure can lower the extinction probability at short times from a repressed preparation, because erasure also removes repressive marks and permits access to the protected active states of this source. It is a counterexample to universal monotonicity in the actuator, which is one reason the amplitude certificate of Theorem 4.1 is stated against arbitrary predictable policies rather than against constant actions.

C Certificates for the N-site source

This appendix states how the certificates of Section 7 were produced and verified, and displays the one that carries the paper's negative conclusion. All vectors, all slacks and the reconstruction of the source matrices are in the data file `site_certificates.json`; the source module is `site_source.py`, the driver that produces the certificates is `make_certificates.py`, and `verify_certificates.py` replays every check listed below. Every displayed inequality is checked with `fractions.Fraction` arithmetic against a source matrix reconstructed from (17), the binomial allocation law and the death rule, never against stored matrix entries. Floating computation is used only to propose candidate vectors.

C.1 Reconstruction and state order

States are generated as (a, r) with a increasing outermost and r increasing innermost, subject to $a + r \le N$, which reproduces (UU, UR, RR, AU, AR, AA) at $N = 2$. The chemical generator uses (17) with rational e; the single-daughter marginal is $\Lambda_{(a,r),(a',r')} = 2^{-(a+r)}\binom{a}{a'}\binom{r}{r'}$; the first-moment matrix is (4) with $L = 2\Lambda$; and the joint daughter law pairs (a', r') with $(a - a', r - r')$ at the same probability, so sibling correlation is retained in Φ and only the mean operator uses Λ.

C.2 Three kinds of certificate

Spectral sign (Table 2, first column). For each rational e the matrix $M = -A_e$ is a Z-matrix, and all its leading principal minors are computed by exact elimination. By (24) their positivity is equivalent to $\mathrm{s}(A_e) < 0$. Bisection on e at denominator 10^6 gives the brackets; the sign at each endpoint is additionally confirmed by a positive rational vector and (16). At $e = 3/10$ the resulting two-sided brackets are

$$\mathrm{s}(A_{.3}) \in \left[-\tfrac{7}{500}, -\tfrac{139}{10000}\right] \ (N = 4), \qquad \mathrm{s}(A_{.3}) \in \left[\tfrac{239}{10000}, \tfrac{3}{125}\right] \ (N = 6),$$

with the 15- and 28-dimensional witnesses recorded in the data file.

Contraction weights (Table 2, last two columns). The Perron vector of $A_{1/2}$ is computed in floating arithmetic, rounded to a positive rational vector w, and the inequalities $A_{1/2}w \le -\gamma w$ are then verified exactly for a rational γ rounded down from $\min_i(-(A_{1/2}w)_i/w_i)$. Only the verified inequality is used; no eigenvalue rounded to a favourable sign enters a proof.

Amplitude floors (Table 3). The required object is a single q with $\Phi_{\text{base}}(q) \le 0$ and $\Phi_{v_{\max}}(q) \le 0$. The minimal such q is the least fixed point of the componentwise maximum

$$\mathcal{M}(q)_i = \max\{\Phi_{\text{base}}(q)_i,\ \Phi_{v_{\max}}(q)_i\},$$

which is reached by monotone iteration from 0 and polished by Newton's method on the active selection. Rounding that fixed point is delicate because the inequalities are tight there, so instead of rounding we retract along the segment of Lemma 3.1: with $h = 1 - q$ and $a = 99/100$, the vector $1 - ah$ satisfies

$$\Phi_v(1-ah) \le a\Phi_v(q) - a(1-a)\, b \odot D(h,h) \le -\tfrac{99}{10000}\, b \odot D(h,h) < 0$$

for every admitted v, with a strictly positive margin that dominates the rounding perturbation. The retracted vector is then written with denominator 10^8 (or 10^{10} where the margin is smaller) and both inequality systems are verified exactly. The cost is a 1% reduction of η, which is included in every value reported in Table 3.

C.3 The six-site amplitude floor

At $N = 6$ with baseline $e = .01$ and eraser amplitude $v_{\max} = .29$, the vector $q = 10^{-8} \cdot Q$ with

$$\begin{aligned} Q = (\ & 94286500,\ 98272968,\ 99056686,\ 99331683,\ 99457438,\ 99525067,\ 99565657,\\ & 70705469,\ 86090037,\ 93106756,\ 96203017,\ 97651735,\ 98389280,\\ & 61949773,\ 75113027,\ 85816369,\ 91719992,\ 94863984,\\ & 58024251,\ 68295718,\ 77595263,\ 86256272,\\ & 55984798,\ 63905275,\ 71672698,\\ & 54795331,\ 60968199,\\ & 54037279\ ) \end{aligned}$$

in the state order $(0,0), (0,1), \ldots, (0,6), (1,0), \ldots, (1,5), \ldots, (6,0)$ satisfies $\Phi_{.01}(q) \le 0$ and $\Phi_{.3}(q) \le 0$ in exact rational arithmetic, with largest coordinate slacks $-6.47 \cdot 10^{-8}$ and $-3.30 \cdot 10^{-6}$ respectively. By Theorem 4.1, every predictable policy with eraser amplitude at most .29 — including cell-specific policies, at unbounded budget and unbounded horizon — leaves the six-site population alive with probability at least $1 - (1-\eta)^n$ from n founders of type $(6,0)$, where $\eta = 1 - 54037279/10^8 = .45962721$.

C.4 Reproduction

The single command `python verify_certificates.py` replays every exact check reported in Section 7: it reconstructs the source at each N, re-verifies the M-matrix brackets, the Collatz–Wielandt witnesses, the contraction weights, the baseline exposure certificates, the amplitude floors and the added-death thresholds, and prints each verdict. The two-site checks of Section 6 and the enclosure (22) are replayed by the scripts inherited from the two-site study. No floating optimizer and no Monte Carlo simulation is needed to check any claim; floating computation appears only in the candidate search and in the clearly labelled diagnostics of Remark 7.2 and Figure 2.

D Evidence and formalization boundary

We distinguish four kinds of evidence and label every quantitative claim by the kind that supports it. These are different kinds of evidence, not interchangeable levels of confidence; in particular, compiled algebra does not formalize the surrounding stochastic theorem.

- **C** Conventional mathematical proof. This covers the controlled construction of Section 2, the segment lemma, the exposure floor (Theorem 3.3), the actuator corollary, the amplitude floor (Theorem 4.1), the time floor (Proposition 4.4), the weighted-drift proposition, the logarithmic order (Corollary 5.2), the dimension-free sufficiency of Proposition 7.3, and the concentration and administered-amount bounds of Section 8.2.
- **E** Exact finite computation or rigorous enclosure. This covers all rational source residuals of Sections 6 and 7, every entry of Tables 2, 3 and 4, the determinant polynomial (25), the Taylor coefficient of Appendix B, and the outward-rounded dyadic enclosure (22).
- **K** Compiled declarations in Lean 4 with mathlib [10, 20], of the stated scope only: the literal six-state source algebra, the supersolution segment identity, the control inequality (19), the improved weight, the RR coefficient calculation, and auxiliary finite barrier and dose rearrangement identities. Strict receipts record the pinned toolchain, source hashes and warnings-as-errors verification.
- **N** Exploratory floating computation. This covers the schedule curves and the full-withdrawal comparison of Figure 2, the large-N values of Remark 7.2 and the open markers of Figure 3, the candidate vectors later verified exactly, and the floating amplitude optimization quoted in Section 7.3.

Neither the main stochastic theorems nor the validated integrator is claimed to be formalized in a proof assistant. That boundary is part of the result and not an omitted assumption hidden in an axiom. The amplitude floors of Table 3 are a natural candidate for a further finite Lean module, since they reduce to two systems of rational polynomial inequalities in a fixed vector; the stochastic step that turns them into Proposition 7.1 is Theorem 4.1 and is conventional.

The accompanying source package contains `main.tex` and the section files, `references.bib`, vector figures, the certificate data files, the generators and checkers named in Appendix C, and a README with canonical commands. The companion source manuscript for the inheritance construction is Companion source manuscript [9].

References

[1] Krishna B. Athreya and Peter E. Ney. *Branching Processes*, volume 196 of *Grundlehren der mathematischen Wissenschaften*. Springer, Berlin, 1972. doi: 10.1007/978-3-642-65371-1.

[2] Nathalie Q. Balaban, Jack Merrin, Remy Chait, Lukasz Kowalik, and Stanislas Leibler. Bacterial persistence as a phenotypic switch. *Science*, 305(5690):1622–1625, 2004. doi: 10.1126/science.1099390.

[3] Abraham Berman and Robert J. Plemmons. *Nonnegative Matrices in the Mathematical Sciences*, volume 9 of *Classics in Applied Mathematics*. SIAM, Philadelphia, 1994. doi: 10.1137/1.9781611971262.

[4] Simone Bruno, Ruth J. Williams, and Domitilla Del Vecchio. Epigenetic cell memory: the gene's inner chromatin modification circuit. *PLoS Comput. Biol.*, 18(4):e1009961, 2022. doi: 10.1371/journal.pcbi.1009961.

[5] Simone Bruno, Felipe A. Campos, Yi Fu, Domitilla Del Vecchio, and Ruth J. Williams. Analysis of singularly perturbed stochastic chemical reaction networks motivated by applications to epigenetic cell memory. *SIAM J. Appl. Dyn. Syst.*, 23(4):2695–2731, 2024. doi: 10.1137/23M1592389.

[6] Nigel J. Burroughs and Mathilde L. C. Leuridan. Optimising the tumour elimination payoff in cancer therapy. *IET Control Theory & Applications*, 18:1621–1637, 2024. doi: 10.1049/cth2.12701.

[7] Tyler Cassidy, Daniel Nichol, Mark Robertson-Tessi, Morgan Craig, and Alexander R. A. Anderson. The role of memory in non-genetic inheritance and its impact on cancer treatment resistance. *PLoS Comput. Biol.*, 17(8):e1009348, 2021. doi: 10.1371/journal.pcbi.1009348.

[8] Julien Claisse. Optimal control of branching diffusion processes: a finite horizon problem. arXiv:1511.06809, 2015. URL https://arxiv.org/abs/1511.06809.

[9] Companion source manuscript. Conservative molecular inheritance and population risk: extinction and variance ordering, approximation limits, and regrowth, 2026. Manuscript supplying the inheritance construction of Section 5; local file `Conservative_Molecular_Inheritance_Population_Risk`. Source hash recorded in the reproduction package.

[10] Leonardo de Moura and Sebastian Ullrich. The Lean 4 theorem prover and programming language. In *Automated Deduction – CADE 28*, volume 12699 of *Lecture Notes in Comput. Sci.*, pages 625–635. Springer, Cham, 2021. doi: 10.1007/978-3-030-79876-5_37.

[11] Ian B. Dodd, Mille A. Micheelsen, Kim Sneppen, and Geneviève Thon. Theoretical analysis of epigenetic cell memory by nucleosome modification. *Cell*, 129(4):813–822, 2007. doi: 10.1016/j.cell.2007.02.053.

[12] Einar Bjarki Gunnarsson, Benedikt Vilji Magnússon, and Jasmine Foo. Optimal dosing of anti-cancer treatment under drug-induced plasticity. arXiv:2412.16391, version 2, 2024. URL https://arxiv.org/abs/2412.16391.

[13] Marek Kimmel and David E. Axelrod. *Branching Processes in Biology*, volume 19 of *Interdisciplinary Applied Mathematics*. Springer, New York, 2 edition, 2015. doi: 10.1007/978-1-4939-1559-0.

[14] Ramon E. Moore, R. Baker Kearfott, and Michael J. Cloud. *Introduction to Interval Analysis*. SIAM, Philadelphia, 2009. doi: 10.1137/1.9780898717716.

[15] N. S. Nedialkov, K. R. Jackson, and G. F. Corliss. Validated solutions of initial value problems for ordinary differential equations. *Appl. Math. Comput.*, 105(1):21–68, 1999. doi: 10.1016/S0096-3003(98)10083-8.

[16] Yaara Oren, Michael Tsabar, Michael S. Cuoco, Liat Amir-Zilberstein, Heidie F. Cabanos, Jan-Christian Hütter, et al. Cycling cancer persister cells arise from lineages with distinct programs. *Nature*, 596: 576–582, 2021. doi: 10.1038/s41586-021-03796-6.

[17] Sydney M. Shaffer, Margaret C. Dunagin, Stefan R. Torborg, Eduardo A. Torre, Benjamin Emert, Clemens Krepler, et al. Rare cell variability and drug-induced reprogramming as a mode of cancer drug resistance. *Nature*, 546:431–435, 2017. doi: 10.1038/nature22794.

[18] Sreenath V. Sharma, Diana Y. Lee, Bihua Li, Margaret P. Quinlan, Fumiyuki Takahashi, Shyamala Maheswaran, et al. A chromatin-mediated reversible drug-tolerant state in cancer cell subpopulations. *Cell*, 141(1):69–80, 2010. doi: 10.1016/j.cell.2010.02.027.

[19] Gustaf Söderlind. The logarithmic norm. history and modern theory. *BIT Numer. Math.*, 46(3):631–652, 2006. doi: 10.1007/s10543-006-0069-9.

[20] The mathlib Community. The Lean mathematical library. In *Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP 2020)*, pages 367–381, New York, 2020. ACM. doi: 10.1145/3372885.3373824.